

EV316932639

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**GENERATION OF CONFIGURATION INSTRUCTIONS
USING AN ABSTRACTION TECHNIQUE**

Inventors:

Matthew E. Miller

Bruce L. Chamberlin

Michael D. Lewis

and

Shaun H. Blackmore

ATTORNEY'S DOCKET NO. MS1-1615US

1 **TECHNICAL FIELD**

2 This invention relates to the configuration of programmable machines, and, in a
3 more particular implementation, to the generation of instructions to configure one or
4 more computers.

5 **BACKGROUND**

6 Mechanisms exist for automatically installing a software product on a computer
7 with a minimum of human intervention. These mechanisms typically provide one or
8 more configuration files that define a series of steps that the computer should perform to
9 install the software product. As the installation proceeds, the computer accesses
10 information from these files. Such information may include a collection of parameters
11 that define values that govern the configuration process. As a result of the configuration,
12 the computer is transformed from a generic processing unit to a specific unit for
13 performing a specific role (such as a web server). In this sense, the installation of
14 operating system programs and application programs can be viewed as “building” a
15 machine.

16 In many business environments, building a computer infrastructure solution may
17 require configuring a collection of computers. There is typically a cooperative
18 relationship between the computers in such a collection. For instance, a particular
19 business environment may include a cluster of computers (defining “nodes”) configured
20 to collectively perform a function. In another common arrangement, a business
21 environment may include one or more backup computer systems that mirror the operation
22 of a normal production computer system. Accordingly, the task of configuring these
23 collections of computers must take into account the environment-specific interaction
24 between these computers.

1 There are many challenges involved in configuring groups of computers. For
2 instance, in implementing a particular solution, a business can select from a very wide
3 variety of computer architectures, machine interconnection arrangements, vendors,
4 models, operating systems, application programs, and so on. A system provider (e.g., a
5 technician or engineer) must therefore be amply familiar with the technical details
6 associated with each of these design options. For this reason, the task of generating
7 configuration instructions often requires the services of highly skilled technicians or
8 engineers who have a significant degree of experience with different setup options. And
9 even with such trained individuals, it may require substantial effort and time to sort out
10 all of the complexities involved in generating instructions for sites that include many
11 different computers.

12 Moreover, because a solution may involve a unique aggregation of design options,
13 a system provider may have difficulty modifying a solution generating for one business
14 site for use in another business site with similar requirements. That is, the unique
15 interdependencies in the solution may prevent the system provider from modifying a
16 prior solution in a modular fashion to provide an updated solution. This difficulty may
17 force the system provider to essentially start from scratch in constructing a new solution
18 for a new business site. Needless to say, these problems may impact the business in a
19 negative way. For instance, these inefficiencies may equate to longer than expected
20 delays and higher than expected costs associated with the configuration task.

21 Tools exist to assist system providers in generating instructions for the
22 configuration of computers. Nevertheless, there is room for substantial improvement in
23 such tools. For instance, the tools may still approach the task of generating instructions
24 on a case by case basis, providing scripts and protocols narrowly tailored to each case.
25 Hence, these tools may not evolve gracefully to accommodate new architectures,

1 interconnection arrangements, operating systems, and application programs introduced to
2 the market. In other words, these tools may lack a general design paradigm that prevents
3 them from efficiently managing the ever-changing myriad of design options available in
4 constructing computer solutions.

5 The above-described lack of a general design paradigm also negatively impacts
6 the efficient transfer of configuration information between different technical personnel.
7 The transfer of information is typically accomplished by specifying an identified
8 computer setup in a burdensome and ad hoc manner, that is, by identifying critical
9 parameters, scripts, and interconnection arrangements using a variety of communication
10 techniques (such as drawings, computer documentation, oral instruction, and so on).
11 Because of the ad hoc nature of this information transfer, faulty or insufficient
12 information may be conveyed, requiring potentially time consuming and costly rework.
13 And even when accurate information is conveyed, the conventional approach does not
14 provide an efficient technique for quickly deploying the solution specified in the
15 instructions. These inefficiencies may grow in proportion to the complexity of the
16 configuration task. As stated above, many sites require the configuration of multiple
17 machines having a cooperative relationship. There is no efficient technique for
18 conveying consolidated instructions that can be used to configure machines in this kind
19 of complex environment.

20 Accordingly, there is a general need in the art for a more efficient and user-
21 friendly technique for generating instructions used to configure programmable machines,
22 such as computers.

23
24
25

1 **SUMMARY**

2 A definition module manipulates scripts used to configure one or more machines
3 using a collection of abstract configuration objects. The configuration objects are drawn
4 from a library of such objects. The configuration objects include a variety of attributes
5 and parameters associated therewith. Pre-established templates define how the
6 information provided in the library is organized to build groups of machines. Some of
7 the objects in the library can include parameters having unspecified values. The
8 definition module uses an inheritance mechanism to automatically resolve these
9 unspecified parameter values on the basis of the context in which the object is being used
10 in a particular build configuration. The manipulation of scripts on an abstract level is a
11 particularly efficient and user-friendly method for building machines. For instance,
12 configuration instructions can be modified on a modular basis without incurring
13 burdensome rework of the entire build configuration. In this manner, the user is insulated
14 from much of the intricacy involved in configuring the machines.

15 The configuration instructions can be consolidated into packages. In one case, a
16 single package can be used to completely configure all of the machines located at a
17 particular site. Further, a site can successively access multiple packages to provide
18 multiple respective configurations to suit different processing environments (e.g., to suit
19 different processing loads). The definition module can efficiently transmit or receive
20 packages to/from other sites. One such other site is a central repository of configuration
21 objects made available to users. The transfer of packages represents an efficient
22 technique for sharing technical information between users, or between a head-end entity
23 and its subscribers.

24 The definition module also includes a graphical user interface having a tree
25 display section for displaying a hierarchical tree of objects associated with the

1 configuration instructions, a parameter display section for displaying parameters
2 associated with the objects in the tree display section, and a properties display section for
3 displaying properties associated with objects in the tree display section or the parameters
4 in the parameter display section. The user interface display provides a convenient
5 mechanism for manipulating the configuration objects, and for conveying information
6 regarding the configuration objects.

7 Additional benefits of the above-described mechanism for generating
8 configuration instructions are set forth in the following discussion.

9

10 **BRIEF DESCRIPTION OF THE DRAWINGS**

11 Fig. 1 shows an exemplary system for defining instructions for the configuration
12 of a collection of machines, and for using those instructions to configure the collection of
13 the machines.

14 Fig. 2 shows an exemplary collection of machine resources that can be built using
15 the system of Fig. 1.

16 Fig. 3 shows a technique for configuring a machine by breaking the configuration
17 task into a series of stages, each of which may include a plurality of phases.

18 Fig. 4 shows information stored in a library provided in the system of Fig. 1.

19 Fig. 5 shows a hierarchy of objects in a particular configuration site, and also
20 demonstrates how parameters associated with a particular object may inherit the values of
21 objects located higher in the hierarchy.

22 Fig. 6 shows an exemplary markup language schema used to specify the
23 organization of information in a template provided by the system of Fig. 1.

24 Fig. 7 shows an exemplary markup language file that conforms to the schema
25 shown in Fig. 6.

1 Fig. 8 demonstrates how imported template information and library information
2 provided in the system of Fig. 1 combine to provide files from which a specific machine
3 can be built.

4 Fig. 9 shows an exemplary Machine.INI file generated by the system of Fig. 1.

5 Figs. 10-33 show different UI presentations provided by the system of Fig. 1 used
6 to perform a variety of tasks.

7 Fig. 34 shows an exemplary procedure for creating and applying configuration
8 files using the system of Fig. 1.

9 Fig. 35 shows an exemplary procedure whereby parameters associated with
10 objects in a configuration hierarchy can inherit values from objects located higher in the
11 hierarchy.

12 Fig. 36 shows an exemplary procedure for building a machine by taking into
13 account interdependencies involved in configuring multiple machines.

14 Fig. 37 shows an exemplary computer architecture that can be used to implement
15 build management logic shown in Fig. 1.

16 The same numbers are used throughout the disclosure and figures to reference
17 like components and features. Series 100 numbers refer to features originally found in
18 Fig. 1, series 200 numbers refer to features originally found in Fig. 2, series 300 numbers
19 refer to features originally found in Fig. 3, and so on.

20

21 **DETAILED DESCRIPTION**

22 This disclosure pertains to the configuration of electronic programmable
23 machines (henceforth referred to as simply “machines”). The task of configuring the
24 machines is also referred to as “building” herein; this terminology reflects the fact that
25 the configuration effectively transforms the machine from a generic machine into a

1 specific machine that serves a specified role. In one exemplary implementation, the
2 techniques described herein can be applied to the configuration of a single electronic
3 machine in either a home-use context or business-use context (or other context). In
4 another exemplary implementation, the techniques described herein can be applied to the
5 configuration of multiple machines in a home-use context or business-use context (or
6 other context). In the case of multiple machines, there may be interrelationships in the
7 functionality provided by the machines that should be taken into account in the
8 configuration of any member of such a collection of machines.

9 The term “machine” used herein has broad connotation. The term generally
10 represents any kind of resource that can be configured by installing machine instructions
11 on the resource. A “machine” is most often identified in this disclosure as some kind of
12 computer, such as a general purpose computer or an application-specific computer. A
13 computer is programmed by storing machine instructions, in its memory, that define its
14 operating system and application programs. Servers represent one class of computers
15 designed to provide services to users via a network, such as the Internet, an intranet, a
16 LAN, etc. In addition, in other implementations, the techniques described herein can be
17 applied to other kinds of devices, such as routers, network switches, network storage
18 devices, programmable control systems, and so on.

19 The term “configuration instructions” used herein has broad connotation as well.
20 It refers to any kind of information that can be generated that specifies the steps that a
21 machine should sequence through to install a particular software product. The
22 configuration instructions may include scripts that specify a series of actions that the
23 machine should perform. The configuration instructions may also include one or more
24 parameters that define environment-specific settings that govern the setup. The
25 configuration instructions may also include one or more links (i.e., references) to other

1 collections of configuration instructions. The configuration instructions may prompt the
2 machine to install an operating system, application program, and/or other kind of code or
3 information that defines its functionality.

4 Generally speaking, the tools described herein for generating instructions for
5 configuring machines can be implemented in hardware, software, or combination of
6 hardware and software.

7 This disclosure is organized as follows. Section A of this disclosure provides an
8 overview of exemplary techniques and implementations for generating and applying
9 configuration instructions. Section B discloses exemplary user interface (UI)
10 presentations that assist a user in the generation of the configuration instructions, as well
11 as other tasks. Section C provides additional exemplary details regarding the operation of
12 a system for generating configuration instructions described herein. And Section D
13 provides an exemplary computer environment for generating the configuration
14 instructions.

15

16 **A. Overview of Techniques and Implementations**

17 *Exemplary System*

18 Fig. 1 shows an exemplary system 100 for generating instructions for the
19 configuration of a collection of machines, and for applying those instructions to the
20 collection of machines. More specifically, Fig. 1 shows multiple sites: site X 102, site Y
21 104, and site n 106 (where “n” generally indicates the last of a plurality of sites). A “site”
22 generally represents infrastructure associated with a collection of machines to be built
23 (i.e., configured). For instance, different sites can correspond to different businesses
24 having respective groups of machines to be configured. Alternatively, different sites can

1 correspond to different parts of an infrastructure associated with a single business. Other
2 variations on this site-based motif are possible.

3 The functionality associated with the exemplary site X 102 is divided into three
4 categories shown in Fig. 1. In a first category, equipment 108 generally represents
5 functionality for defining or generating the configuration instructions, as well as other
6 high-level managerial tasks, such as the transfer or receipt of configuration instructions
7 to/from other sites. In a second category, equipment 110 represents functionality for
8 coordinating the actual building of the machines. In a third category, equipment 112
9 represents the machines to be built. In this last category of equipment 112, three server-
10 type machines are shown, although any number or type of machines can be built
11 (including only one machine). The functionality in each of the above-described
12 categories will be discussed in turn in the following explanation.

13 To begin with, equipment 108 includes a configuration definition and
14 management module 114 (referred to for brevity as the “definition module” 114
15 henceforth). The definition module 114 generates the instructions used to build the
16 machines in the collection of equipment 112. To perform this task, the definition module
17 114 includes build management logic 116 and automatic purposing logic 118. The
18 automatic purposing logic 118 provides functionality for manipulating a collection of
19 scripts and associated parameters that can be assembled to create the configuration
20 instructions. An exemplary software product that can be used for this task is the
21 Automatic Purposing Framework (APF) software product produced by Microsoft
22 Corporation of Redmond, Washington. The build management logic 116 provides a so-
23 called “front end” to the automatic purposing logic 118. More specifically, the build
24 management logic 116 provides a technique for abstracting parameters provided by the
25 automatic purposing logic 118 so that these parameters can be handled in a more efficient

1 and flexible manner than heretofore provided. In other words, the build management
2 logic 116 “manages” the parameters provided by the automatic purposing logic 118 so
3 that they can be manipulated in more effective ways than heretofore provided. The build
4 management logic 116 also provides functionality for providing user interface (UI)
5 presentations that allow a user to interact with the build management logic 116 in a user-
6 friendly and efficient manner. The build management logic 116 also performs other
7 tasks, to be discussed in due course. The build management logic 116 is also referred to
8 by an exemplary name of “APF Explorer” software product. However, it should be noted
9 that the design concepts discussed here are not limited in application to the Automatic
10 Purposing Framework software product.

11 The build management logic 116 and the automatic purposing logic 118 interact
12 with a database 120 and a template storage 122. The database 120 stores a collection of
13 objects that allow a user to manage the parameter information provided by the automatic
14 purposing logic 118. More specifically, the database 120 can be conceptualized as
15 including at least three groupings of information. In a first grouping, an APF library 124
16 (henceforth simply “library 124”) defines building blocks that a user can string together
17 to provide configuration instructions for a specific collection of machines (or other
18 resources). More specifically, as will be discussed in connection with Fig. 5, such
19 building blocks are grouped into different categories, such as objects pertaining to
20 different available vendors, objects pertaining to different operating systems, objects
21 pertaining to different machine functions, objects related to stages and phases involved in
22 configuring machines, and so on. The objects stored in the library 124 are maintained in
23 a relatively general, generic, and standardized form. For instance, while some parameters
24 associated with the objects in the library 124 are specified in advance, other parameters
25 are left open-ended and will assume specific values only when used to construct

1 configuration instructions for a specific combination of machines at a site. For example,
2 a parameter in the library 124 might have a generic value of “Default.DistUser,” which
3 might resolve to the specific value of “administrator” when used in a particular build
4 environment. The generic and flexible nature of the library 124 has notable advantages,
5 as it allows common building blocks to represent many different configuration options;
6 this is in contrast to prior solutions, which attempted to construct unique proprietary
7 solutions for different configuration options without regard to any common principles
8 binding different solutions together.

9 A second grouping in the database 120 pertains to a site tree arrangement of
10 objects 126 (henceforth simply “site tree 126”). More specifically, the building blocks
11 provided in the library 124 can be assembled to form a hierarchical organization of
12 objects representative of machines to be configured at a particular site (such as site X
13 102). For example, assume that a particular site includes two groups of machines, G1
14 and G2, where group G1 includes machines M1 and M2, and group G2 includes
15 machines M3 and M4. In this case, the site tree 126 would represent a collection of
16 objects taken from the library 124 that represent such an organization of machines.

17 A third grouping in the database 120 represents packages 128. Packages 128 refer
18 to collections of objects in the database 120 that are assembled together as a single file.
19 This allows the packages to easily be transferred from one site to another.

20 The database 120 can be implemented using different kinds of technology, such
21 as structured query language (SQL) technology.

22 On the other hand, the template storage 122 provides information pertaining to
23 templates. Templates serve at least two roles in the definition module 114. According to
24 a first role, the templates can be used to provide a skeleton representation of a machine or
25 group of machines specified in the site tree 126 (or elsewhere). This skeleton

1 representation conforms to a predetermined form (schema) and contains enough
2 information, in combination with information in the library 124, to reconstruct a
3 description of a group of machines or a single machine. In one exemplary
4 implementation, the template information is expressed in a markup language, such as the
5 Extensible Markup Language (XML). Templates stored in the template storage 122 can
6 be exported and transferred to other sites, where they can be used to generate
7 configuration instructions at such other sites. The build management logic 116 can also
8 cull a subset of information from the database 120 and organize it into a package
9 (represented by package objects 128), and then transfer this package to another site. Fig.
10 shows the transfer of an exemplary template or package 130 to a remote site, such as
11 site Y 104. The transfer of XML templates and packages 130 between sites provides an
12 efficient way of conveying configuration instructions between sites. As mentioned in the
13 BACKGROUND section of this disclosure, the transfer of instructions in prior systems
14 was handled in a non-systematic, not-standardized, ad hoc manner.

15 According to another role, in one implementation, groups of machines in the site
16 tree 126 are specified by filling out a template that governs the organization and nature of
17 machines in the group. The makeup of individual machines in the group is also specified
18 by filling out a template that governs the basic features of the machine. As will be
19 described in greater detail in Section B, this functionality can be implemented by
20 sequencing the user through a series of configuration guidance displays that prompt the
21 user to fill out the predetermined templates.

22 Input devices 132 are provided that allow a user to interact with the build
23 management logic 116. Such input devices 132 may include a conventional keyboard
24 and mouse-type pointing device, but can include any kind of input device (such as voice
25 recognition, touch-sensitive screen input, etc.).

1 A display device 134 is provided for receiving and displaying information
2 generated by the build management logic 116. The display device 134 may comprise a
3 computer monitor or other electronic display apparatus having a graphical user interface
4 (GUI) presentation 136. In one exemplary implementation, the GUI 136 organizes
5 information into three panels or panes. Pane 138 shows the organization of a collection
6 of objects stored in the library 124 of the database 120 in hierarchical tree form. Pane
7 138 also shows the organization of a collection of objects in the site tree 126 of the
8 database 120 in hierarchical tree form. Pane 138 also shows the organization of a
9 collection of objects in the packages 128 of the database 120 in hierarchical tree form.
10 On the other hand, pane 140 shows one or more parameters associated with a highlighted
11 object in pane 138. Pane 142 shows one or more properties associated with a highlighted
12 item in either pane 138 or pane 140. For instance, if a user selects and highlights one of
13 the parameters shown in pane 140, then pane 142 will present information regarding the
14 properties of the highlighted parameter. Additional details regarding the information
15 presented in panes 138, 140, and 142 are provided in Section B below.

16 The build management logic 116 generates one or more files that specify the
17 configuration instructions. A collection of such files is referred to as manifest files 144.
18 The manifest files 144 can include a file having the extension “.INI.” An .INI file
19 specifies the information required to configure a machine. For instance, this file may
20 specify various scripts used to configure the machine, as well as various parameters
21 associated with the scripts. For example, a file “Machine_A.INI” would provide the
22 information required to configure a machine named “Machine_A.” The manifest files
23 144 may also include a so-called answer file that specifies a number of parameters used
24 in installing an operating system on a machine. The manifest files 144 may also include
25 a synchronization file that specifies various dependencies that govern building several

1 machines in parallel. For instance, a particular machine may reach a point in its
2 configuration where it needs the support of (or information provided by) another machine
3 that is concurrently being built. The synchronization file in the manifest files 144
4 provides information that specifies such dependencies.

5 Advancing now to the middle tier of Fig. 1, equipment 110 includes a deployment
6 module 146. The deployment module 146 can represent a server or other kind of
7 computer device that coordinates the configuration of the machines in the collection of
8 equipment 112. It performs this task using build logic 148, in conjunction with the
9 manifest files 144 specified by the build management logic 116. Build logic 148 is also
10 coupled to master dependency logic 150. The master dependency log 150 stores
11 information regarding the progress of each machine in executing its allocated
12 configuration tasks. This information serves as a central repository that can be accessed
13 by any machine in the collection of equipment 112. As will be discussed, the machines
14 use this information to ensure that they perform their allocated configuration tasks in a
15 synchronized and cooperative manner.

16 Finally, in the lowest tier in Fig. 1, equipment 112 is shown as including three
17 machines (152, 154, 156) to be built, such as three servers. The machines (152, 154, 156)
18 can represent standalone units. In this case, these machines (152, 154, 156) can be
19 individually configured without taking into consideration the configuration of other
20 machines. Alternatively, the machines (152, 154, and 156) can be arranged in a group in
21 which the machines (152, 154, 156) interact with each other to collectively provide some
22 functionality. In this case, the machines (152, 154, 156) may have to be configured in a
23 specific order. The synchronization file specified in the manifest files 144 provides
24 information regarding the interdependencies involved in configuring multiple machines.
25 That is, this file ensures that, in the configuration process, no machine attempts to use a

1 resource before it is built. Also, as mentioned above, the master dependency log 150 in
2 the deployment module 146 can be used to inform each machine of the progress of other
3 machines in performing their respective configuration tasks. Thus, as described above,
4 the synchronization file and the master dependency log 150 collectively provide a
5 mechanism by which the machines (152, 154, 156) can coordinate their allocated
6 configuration tasks with respect to each other.

7 Jumping ahead briefly in the sequence of figures, Fig. 2 provides an exemplary
8 grouping of machine resources available at a particular site, as represented in site tree 126
9 of the database 120. The grouping is arranged in tree form to represent the hierarchical
10 relationship between different features of the grouping. A topmost “site Z” object 202
11 generally represents an aggregate of all of the equipment to be configured at a site Z.
12 Objects 204 and 206 represent subgroupings of equipment within the site Z object 202.
13 Namely, object 204 pertains to a first web hosting test lab, while object 206 represents a
14 second web hosting test lab. The first web hosting test lab includes first and second web
15 servers, represented by objects 208 and 210, respectively. The second web hosting test
16 lab includes a generic cluster represented by object 212. The generic cluster, in turn,
17 includes an exemplary collection of four nodes represented by objects 214, 216, 218, and
18 220. In summary, in Fig. 2, the bottom leafs of the tree (corresponding to objects 208,
19 210, 214, 216, 218, and 220) represent actual machines (e.g., computers) that require
20 configuration. The higher nodes in the tree conceptually represent different groupings of
21 the actual machines in the lower leafs of the tree. Again, however, the specific grouping
22 of machines in Fig. 2 is merely illustrative of one of a myriad of different configuration
23 options that the system 100 shown in Fig. 1 can be applied to. Other multi-computer sites
24 may include various groupings of production and backup computers, etc. Further, the
25 system can be used to configure other kinds of machines, such as routers, networks

1 switches, network storage devices, and so on. The definition module 114 can send these
2 kinds of devices a sequence of commands and parameter values to configure them.

3 Returning to Fig. 1, machine 152 illustrates exemplary internal features of one of
4 the machines in the equipment 112. As shown in this figure, machine 152 includes
5 configuration logic 158 that coordinates the configuration of machine A. In one
6 implementation, the manifest files 144 can be stored in the deployment server 146 during
7 initial stages of the initialization process, but can thereafter be downloaded into
8 individual respective machines. Configuration logic 158 can store information obtained
9 from the manifest files 144, as well as additional configuration-related information that is
10 downloaded to machine 152 (or obtained from some other source).

11 The actual “target” of the configuration procedure is the resources 160. The
12 resources 160 can represent memory that is to receive an operating system, one or more
13 application programs, or other information. The resources 160 can also represent other
14 mechanisms that can be changed from one state to another depending on the
15 configuration instructions specified in the manifest files 144. For instance, such other
16 mechanisms could therefore encompass any kind of switching or routing mechanisms.

17 Machine 152 also includes its own progress log 162. The progress log 162 stores
18 an indication of how far the machine 144 has advanced in its configuration. Information
19 stored in progress log 162 can be forwarded to the master dependency log 150 in the
20 deployment module 146; this allows other machines to be apprised of this machine 152’s
21 configuration progress.

22 Generally, Fig. 1 separates different functional aspects of the system 100 into
23 separate modules to facilitate discussion. However, a collection of these modules can be
24 implemented by a single computer device. For instance, the build management logic 116
25 can be implemented on the same computer device as the deployment module 146. On the

1 other hand, other modules that are shown as coupled together can alternatively be
2 implemented as separate computer devices. For instance, the build management logic
3 116 and automatic purposing logic 118 can be implemented on separate computers.
4 Other such implementation arrangements are possible. Generally, any module can be
5 coupled to any other module by a hardwired link, by a network, or by another kind of
6 coupling strategy. As mentioned above, any of the logic shown in Fig. 1 can be
7 implemented as software-based machine instructions, firmware, or a combination of
8 software and firmware.

9 Finally, although not shown, the build management logic 116 can be incorporated
10 within other kinds of application programs and associated business environments, such as
11 order entry programs, image deployment, asset management solutions, etc. In this
12 context, the build management logic 116 is used as a tool within such other kinds of
13 business environments.

14 In one exemplary configuration, each of the machines shown in Fig. 1 is
15 configured by progressing through a series of stages. And each stage, in turn, may
16 include multiple phases that are executed in a prescribed sequence. Fig. 3 shows details
17 of one such strategy for configuring a machine using a series of stages and phases.
18 Namely, the configuration proceeds by sequencing through, in order, a PreOS stage 302,
19 an InstOS stage 304, a CmdLines stage 306, and a Post OS stage 308. The PreOS stage
20 302 includes a series of phases 310 performed in the top to bottom order specified in Fig.
21 3. The PostOS stage 308 includes another series of phases 312 shown in Fig. 3. A first
22 phase list (not shown) can be used to reference the collection of phases 310, and a second
23 phase list (not shown) can be used to reference the collection of phases 312.

24 The PreOS stage 302 involves the configuration of hardware BIOS, RAID arrays,
25 and similar resources. This stage 302 also creates and formats disk partitions for

1 installing an operating system. The InstOS stage 304 initiates the automated installation
2 of the operating system. The CmdLines stage 306 updates the operating system to
3 automatically run the PostOS stage 308 after the installation of the operating system
4 completes and the computer device reboots; hence, the CmdLines stage 308 is performed
5 relatively late in the installation of the operation system. Finally, the PostOS stage 308
6 performs actions that configure the computer for its intended purpose after the operating
7 system installation.

8 Fig. 3 reflects only one exemplary implementation that associates phase lists with
9 the PreOS stage 302 and the PostOS stage 308. However, other implementations can
10 associate phase lists with any other combination of stages, including, in one case, stages
11 302, 304, and 308, and, in another case, all of the stages. Further, the stages shown in
12 Fig. 3 are also exemplary. Other applications may employ different types of stages and
13 different numbers of stages.

14 Illustrations of exemplary phases will be presented in Section B of this disclosure.
15 Generally, a phase invokes one or more scripts that perform a prescribed configuration
16 task. The phase may also specify one or more parameters which supply values used in
17 performing the configuration task. For instance, a parameter might specify a numerical
18 value that governs the setup of a piece of hardware used in the machine. Another
19 parameter might act as a pointer to information stored in a database or on a network (such
20 as the Internet). These are merely illustrations of the flexible role of parameters in the
21 configuration process. Generally, the database 120 stores these parameters, and the build
22 management logic 116 manages these parameters by manipulating these parameters in
23 useful and efficient ways.

24
25

1 *Transfer of Configuration Information*

2 Returning to Fig. 1, a network 164 couples the multiple sites (102, 104, . . . 106)
3 together. The network 164 can represent any kind of network, such as a wide area
4 TCP/IP network (e.g., the Internet), an enterprise-wide intranet, a LAN, and so on. The
5 network 164 can be implemented using various types of hardwired links and/or wireless
6 links. In one configuration, the sites (102, 104, . . . 106) can directly transfer templates
7 and packages to each other using any kind of file transfer technology, such as file
8 sharing.

9 In addition, a head-end site 166 provides configuration-related services to each of
10 the sites (102, 104, . . . 106). The head-end site 166 can be implemented as a server. It
11 can include head-end logic 168 for performing host-related tasks, as well as a database
12 170. Any of the sites (102, 104, . . . 106) can download templates and packages from the
13 database 170 of the head-end site 166 (and can also upload templates and package to the
14 database 170). The head-end site 166 can also coordinate the transfer of templates and
15 packages between sites (102, 104, . . . 106). In operation, URL (Uniform Resource
16 Locator) information (or other identifying information) can be associated with templates
17 and packages stored at the head-end site 166. To retrieve a particular resource, a user can
18 specify the URL corresponding to the desired resource. The head-end site 166 can
19 provide various web pages (not shown) which assist the user in retrieving the desired
20 resource. The web pages can provide search functionality for finding a particular
21 resource based on specified key terms. The web pages can also provide catalog-type
22 listings of available resources. As will be described below, different users can be
23 assigned to different groups. In this case, the head-end site 166 can be configured to
24 provide different levels of access to different groups of users.

1 In one case, a package can be created to configure all of the machines located at a
2 particular site. Thus, it is easy to build a relatively complexity group of interrelated
3 machines by downloading a package from another site or the head-end site 166. In one
4 case, a site can receive multiple packages. These packages specify different
5 configurations for the site's machines to suit different processing objectives and
6 requirements. A site user can thus dynamically change the configuration of the machines
7 located at the site simply by loading another package and using this package to generate
8 new configuration instructions for dissemination to the machines. For example, a site
9 user may choose to dynamically vary the configuration at a site in response to changes in
10 traffic load experienced at the site, or based on a predetermined schedule for such
11 changes. In another implementation, a site can be configured to automatically apply a
12 different package in response to detected changes in the processing environment.
13 Alternatively, a user at the head-end site can configure the machines located at a remote
14 site without involvement of personnel at that remote site, or with minimal involvement of
15 personnel at that site. Still alternatively, the head-end site 166 can be configured to
16 automatically configure machines located at a remote site in response to detected changes
17 in the processing environment affecting that site. Still other permutations on this design
18 strategy are possible.

19 The above-described system can be used in a number of different commercial
20 environments. In one case, a group of users can cooperatively maintain the head-end site
21 166 for their collective use. In this context, users at different sites can upload packages
22 that might be of use to other users, as well as download packages from the head-end 166
23 created by other users. In another case, a third party entity can offer the resources of the
24 head-end site 166 to subscribing users. Such an entity may offer the resources free of
25 charge or for a fee, or on some other basis. For instance, this would provide a convenient

1 mechanism for both computer manufacturers and software developers to disseminate
2 software to interested recipients. Such software can provide the base-level initial
3 programming of the kind illustrated in Fig. 3, as well as application programs or other
4 instructions to already configured machines.

5 Although only one head-end site 166 is shown, the system 100 can include
6 multiple head-end sites.

7

8 *The SQL Database*

9 The object-oriented design paradigm is used to organize information in the
10 database 120. Generally, in this design approach, entities are represented as objects.
11 Objects are organized in hierarchies, where each object may have a child object, and each
12 child object, in turn, may have its own child object. Each object may also have attributes
13 and parameters associated therewith.

14 More specifically, within an exemplary hierarchy, the object-oriented approach
15 provides the following general root objects in the present case: Site (represented the root
16 site); Vendors (representing a collection of vendors); OperatingSystems (representing a
17 collection of top level operating systems); Machinefunctions (representing a collection of
18 top level machine functions); ‘Stage Name’ (representing names of stages, including
19 Boot, Common, PreOSStage, InstOSStage, and PostOSStage); ‘OperatingSystem’
20 (representing each top level operating system); ‘MachineFunciton’ (representing each top
21 level machine function); and ‘Vendor’ representing each top level vendor.

22 Within the context of a particular machine, the following root objects are also
23 available: Machine (representing the current machine’s object); Vendor (representing the
24 current machine’s vendor object); Model (representing the current machine’s model
25 object); OperatingSystem (representing the current machine’s operating system object);

1 MachineFunction (representing the current machine's machine function object); Group
2 (representing the current machine's parent group object); and Default (allowing the
3 referencing of a parameter specified at any level of the machine group or site hierarchy).

4 Further, each object in the database 120 has inherent attributes associated with it.
5 The following exemplary attributes are common to each object: ID (representing an
6 internal ID assigned to the object); Name (representing a name assigned to the object);
7 Description (representing a description of the object); Version (representing the object's
8 version number); Owner (representing an individual authorized to perform some action
9 on the object, such as updating the object); OwnerID (representing an ID assigned to the
10 object's owner); Parent (representing the object's parent); and ParentMachine
11 (Representing the parent machine of the object if it exists).

12 Also, an object may have one or more child objects associated therewith. For
13 instance, the Site object has a MachineGroups child object associated therewith. The
14 MachineGroup object has MachineGroups and Machine child objects associated
15 therewith. The Vendor object has a Models child object associated therewith. The
16 Model object has a Models child object associated therewith. The OperatingSystem
17 object has an OperatingSystems child object associated therewith. The MachineFunction
18 object has a MachineFunctions child object associated therewith. And the Machine
19 object has a Stages child object associated therewith.

20 The database 120 therefore contains fine-grained information regarding items
21 involved in the configuration (where such "items" can refer to objects representing a
22 group of machines, a machine, a stage of a machine's configuration, a phase of a stage, a
23 parameter, and so on). In other words, the configuration items may be thought of as
24 forming a hierarchical tree having a plurality of nodes. Different kinds of information
25 can be "attached" to (i.e., associated with) different nodes in the tree in a highly granular

1 fashion. Some of the fine-grained information simply identifies or describes the
2 configuration items. Some of the information refers to validation rules applicable to the
3 configuration item. (Validation rules govern the kind of error checking that the build
4 management logic 116 performs when a user enters information regarding any of the
5 configuration items. For instance, the validation rules may ensure that the entered data
6 has a pre-specified format or type, etc.) Some of the information refers to how the
7 configuration items are displayed on the user interface 136 or printed out in the manifest
8 files 144, or refers to the behavior of the thus displayed configuration items (e.g., how the
9 user can subsequently edit the configuration items). Some of the information pertains to
10 the individual who owns the configuration item. An “owner” may correspond to the
11 individual who created or specified the parameter; therefore, that owner alone may have
12 access rights to change it or export it, etc. Some of the information pertains to the
13 version of the objects and parameters stored in the database 120. In one implementation,
14 a new version of an object or parameter is created each time it is modified. This is
15 merely a representative sampling of different kinds of information that can be associated
16 with the “nodes” in the trees provided in the database 120. It bears repeating that the
17 information stored in the database 120 mainly pertains to pointer type information which
18 ultimately resolves by forming a series of instructions used to configure a collection of
19 machines; however, in interacting with the database 120, the user is not typically
20 manipulating the scripts themselves, but an abstract object framework that represents the
21 scripts, and ultimately refers to the scripts.

22 Fig. 4 shows an exemplary organization of information stored in the library 124 in
23 greater detail, and is used as a vehicle for explaining additional features and merits of the
24 database 120 organization. The library 124 contains vendor information 402 that
25 identifies different manufacturers (or, more generally, providers) of machines and other

1 configurable resources. The library 124 also contains operating system information 404
2 that pertains to a collection of operating systems that can be installed on the machines as
3 part of their configuration. The library 124 also includes machine function information
4 406 that specifies the functions performed by the machines. For instance, a machine can
5 be configured such that it serves the role of a web server; hence, a web server is one of
6 the functions that this machine can perform if properly configured. The library 124 also
7 includes stage information 408, phase information 410, and phase list information 412.
8 This information specifies stages, phase lists, and individual phases that can be used to
9 configure a machine at a particular site. The library 124 can also provide parameter type
10 information 414 that identifies a variety of types associated with parameters specified in
11 the library 124. The library 124 can also include script information 416 that provides the
12 actual instructions used to perform a particular configuration task within a configuration.
13 For instance, different phases can include pointers to scripts which respectively perform
14 the tasks associated with the phases.

15 The site tree 126, on the other hand, typically contains information regarding the
16 site that provides the machines to be configured. The site tree 126 contains group
17 information regarding an organization of machines in one or more groups. The site tree
18 126 also includes machine information regarding machines provided in the individual
19 groups, each of which includes stage, phase list, and phase information associated
20 therewith as specified in the library 124. As mentioned above, the above-identified
21 information in the library 124 and the site tree 126 can be associated with one or more
22 objects in the database 120.

23 Different parameters pertain to different respective aspects of the configuration
24 operation. Accordingly, different parameters can be associated with different respective
25 objects in the database 120. For example, a collection of parameters may have a global

1 bearing on the configuration of multiple machines within a site, while another collection
2 of parameters have a relatively local role in executing one phase within a collection of
3 phases that make up a stage. In the former case, the database 120 can associate the high-
4 level parameters with corresponding high-level objects within the database 120. In the
5 latter case, the database 120 can associate the low-level parameters with corresponding
6 low-level objects within the database 120. Fig. 4 illustrates the above-described
7 object/parameter mapping as follows: parameter information 418 is associated with
8 vendor information 402; parameter information 420 is associated with operating system
9 information 404; parameter information 422 is associated with machine function
10 information 406; parameter information 424 is associated with stage information 408;
11 parameter information 426 is associated with phase information 410; and parameter
12 information 428 is associated with phase list information 412. The information shown in
13 Fig. 4 is merely illustrative of the kinds of information that can be stored in the database
14 120; other implementations can include additional configuration objects, or can omit
15 certain configuration objects shown in Fig. 4. Parameters can be associated with
16 different nodes in the site tree 126 in a similar manner to that described above.

17 Fig. 4 also conceptually shows an excerpt from the phase information 410 and
18 associated parameter information 426 of the database 120. The excerpt shows a set of
19 phases 430 that compose a phase list used in a particular stage in the configuration of a
20 machine. As indicated in this figure, each phase can have one or more parameters
21 associated therewith. For example, in an illustrative case, phase 432 can have parameters
22 434, 436, 438, and 440 associated therewith.

23 As described above, the library 124 does not attempt to capture the complete
24 universe of parameter value permutations possible in different applications. Rather, the
25 library 124 specifies fixed values for a first class of parameters, but the library 124 leaves

1 the values of other parameters open-ended (that is, unspecified) for a second class of
2 parameters. The values for the parameters in the second class can be provided (or
3 instantiated) when objects in the library are assembled together to define a particular
4 configuration environment involving a collection of machines in the site tree 126. The
5 values for the first class of fixed parameters can be specified in advance because, for
6 instance, these values do not change depending on the design context in which their
7 associated configuration object is used. Fig. 4 conveys these concepts in high-level form
8 using the particular collection of parameters (434, 436, 438, and 440) associated with
9 phase 434. Namely, parameters 434 and 440 have fixed values that do not change when
10 the configuration object representing the phase 434 is “inserted” into a specific design
11 solution in the site tree 126. On the other hand, parameters 436 and 438 have open-ended
12 (unspecified) values. The values for these parameters can be specified when the
13 configuration object representing the phase 432 is inserted in a specific design solution.

14 Fig. 5 shows how the open-ended values associated with the phase 432 in Fig. 4
15 are assigned specific values when placed in the context of a specific design solution 500
16 in the site tree 126. Namely, the design solution 500 is provided at a site W. Site W is
17 represented by site object 502 in a hierarchy associated with the design solution 500 in
18 the site tree 126. The site W includes a collection of different groups of machines. One
19 of the groups is represented in the design solution 500 as group object 504, and one of the
20 machines in one of the groups is represented by machine object 506. A plurality of
21 stages, phase lists, and phases are associated with each machine to be configured. Stage
22 object 508 represents one of the stages in the design solution 500, phase list object 510
23 represents one of the phase lists in the design solution 500, and phase 432 represents one
24 of the phases in the design solution 500. Parameter 440 in phase 432 can point to a script
25 512 that performs the tasks allocated to the phase 432.

1 Section B explains how the build management logic 116 can create the type of
2 design solution 500 shown in Fig. 5. By way of preview, the collection of objects can be
3 created from “scratch” by defining the salient features of the objects and their
4 interrelations. This can be performed using one or more configuration guidance
5 presentations that sequence the user through a series of panels designed to solicit
6 information from the user regarding the properties of the objects in the design solution
7 500. In this manner, the properties in the solution are forced to conform to a pre-
8 established template. Alternatively, a user can create a design solution for a particular
9 site by importing such design solution from an external source (e.g., from an external
10 site). Alternatively, a user can create a design solution by modifying a pre-existing
11 design solution. The build management logic 116 also permits users to update the
12 information stored in the library 124, such that the database 118 continues to support a
13 wide variety of configuration object types.

14 Assume, as explained above, that the phase 432 is “inserted” within the design
15 solution 500. Parameters 434 and 440 have fixed values that do not change when the
16 object 432 is inserted in the design solution tree 500. For instance, these parameters
17 might specify that computer resources are mapped to specific fixed drives. On the other
18 hand, as noted before, parameters 436 and 438 have open-ended values. In the
19 illustrative case of Fig. 5, the value of parameter 436 is defined within object 508 (on the
20 stage level), and the value of parameter 438 is defined within object 504 (on the group
21 level). Generally, to resolve the value of a variable parameter, the build management
22 logic 116 steps up the design solution tree 500 to determine whether each successively
23 higher object contains information which resolves the value of the parameter. Line 514
24 in Fig. 5 represents the resolution of the value of parameter 436, while line 516 represents
25 the resolution of the value of parameter 438.

1 The strategy depicted in Fig. 5 therefore provides a flexible technique for
2 specifying parameter values within a specific design solution. Whereas heretofore
3 technical personnel specified the collection of parameters used in a design solution on a
4 case by case basis, the design paradigm shown in Fig. 5 provides general rules for
5 inferring the values of many parameters, so that the specific values are automatically
6 assigned to objects based on the context of their use within a tree of other objects.
7 Hence, the design paradigm shown in Fig. 5 has the potential for defining configuration
8 instructions in a much more user-friendly and flexible manner than heretofore provided.
9 The relative ease at which configuration instructions can be created by the above-
10 described strategy also allows solution providers having relatively low levels of
11 experience to generate configuration instructions.

12 Fig. 5 also illustrates the ability to override parameter values set at higher levels
13 in the hierarchical tree. For example, machine object 506 specifies a parameter value P_d .
14 However, phase 432 also specifies a value for P_d via parameter 434. The value specified
15 locally within phase 432 takes precedent over the more remote value specified at the
16 machine level. In general, the design solution 500 shown in Fig. 5 can be viewed as a
17 large linked list. A value of a parameter associated with an object in the list is defined by
18 the closest object higher in the hierarchy which specifies such value.

19 The build management logic 116 can use different techniques to perform the
20 parameter referencing functions described above. As noted above, the library 124 can
21 represent different configuration items using an object-oriented design approach in which
22 individual configuration items are represented as objects having properties and various
23 relationships to other objects (such as child objects) in a hierarchy of objects. Values in
24 the hierarchy can be specified using so-called expressions. For example, assume that an
25 object named "Machine" has a property called "Name." In this case, the expression

1 “=Machine.Name” will return the name of the Machine object. The conventional “.”
2 notation is used to string together a sequence of linked elements in the hierarchy.

3 For example, the following expressions return the following results in the context
4 of the exemplary hierarchical tree shown in Fig. 2.

5 a) The expression “=Site.Web Host Test Lab 1” returns the result:
6 Object<Machine Group>

7 b) The expression “Site.Web Hosting Test Lab1.Web Server 1” returns the result:
8 Object<Machine>.

9 c) The expression “Site.Web Hosting Test Lab 2.Generic Cluster.Node 1” returns
10 the result: Object<Machine>.

11 d) The expression “Site.Web Hosting Test Lab 1.Web Server 1.Common” returns
12 the result: Object<Stage>.

13 Parameters associated with a particular object can be assessed using the “.”
14 notation in a similar manner to that described above. For example, assume that the
15 parameter OrgName is assigned the value XYZ_Corp. In this case, the expression
16 “=object.OrgName” would resolve to the specific value “XYZ_Corp.” In addition,
17 expressions can be embedded within parameter values using the following exemplary
18 notation: %<expression>%. For example, assume that a currently identified machine
19 corresponds to Node 1 shown in Fig. 2. In this case, the following expressions return the
20 following results:

21 a) The expression “%<Machine>.Name%.INI” returns the results Node 1.INI.
22 b) The expression “Machine Name is %<Machine>.Name%” returns the result
23 “Machine Name is Node 1.”
24 c) The expression “%<Machine>.IniFileName%” returns the results Node 1.INI.

1 The above-described referencing can be used to provide inheritance in the
2 following manner. For instance, suppose that a phase list in the library was generically
3 identified as “OperatingSystem.PostOSPhaseList.” This generic value would prompt the
4 build management logic 116 to locate the PostOSPhaseList that is associated with the
5 OperatingSystem object. In defining a machine, a user may be able to select among a
6 plurality of operating systems, and hence, a plurality of OperatingSystem objects. Hence,
7 the actual phase list returned by this expression will vary depending on the constitution of
8 the site tree 126. However, there may be many OperatingSystem objects associated with
9 different machines in a vast group of machines. Thus, resolving the precise value of the
10 expression will require successively stepping up the site tree 126 to find the closest object
11 which resolves the value of the parameter. That is, the site tree 126 can be viewed as a
12 large linked list. In evaluating an expression, the build management logic 116 starts at a
13 certain point in this linked list, and subsequently moves up the tree. The build
14 management logic 116 attempts to find the “closest” object up the tree which resolves the
15 unspecified parameter.

16 Another feature of the SQL database 120 is a role-based access mechanism.
17 According to this feature, the database 120 grants different editing/access rights to
18 different respective user groups (associated, in turn, with different roles). For example,
19 the database 120 can grant a library editor full access rights to access, modify, and
20 transfer any information stored in the database 120. However, the database 120 may
21 allow other individuals (such as technicians) permission to only generate .INI files for
22 certain groups of machines, or a particular machine. The particular allocation of access
23 rights can be defined to suit the requirements of a particular business and technical
24 environment. The database 120 can determine what access rights are appropriate to a
25 user by mapping the user’s password to predefined access information associated with the

1 password. This access information defines what the user can “see,” and what the user
2 can subsequently “do” to the information (such as editing the information, exporting the
3 information, generating .INI files, and so on).

4

5 *Templates, Packages, and .INI Files*

6 As explained above, a template provides a skeleton or blueprint which defines the
7 basic organization and constitution of groups of machines and individual machines.
8 Templates serve two roles. First, templates serve as a standardized framework for
9 passing information from one user to another. Such a template is formed by culling all of
10 the values that were specifically set for a machine or group of machines and storing those
11 exported values in the template. The template has a fixed organization of elements.
12 Other information that is not specific to the machine or group of machines can be
13 reconstructed by the recipient from their own library 124. For instance, a template may
14 generally include a pointer to a phase list, but may omit a description of the phases within
15 the phase list. The receiving site can reconstitute the phases based on the local library
16 124 at that site. Second, templates can be used as a blueprint for constructing new groups
17 of machines and new machines. More specifically, a series of configuration guidance
18 panels can be used to construct the new groups and machines based on the templates.

19 Templates are represented in a markup language, such as XML. By way of
20 background, XML is a subset of the Standard Generalized Markup Language (SGML)
21 that enables users to create customized tags that describe the meaning of data, as opposed
22 to the presentation of data. Generally, an XML document is composed of XML objects,
23 each of which includes a start tag (such as <author>), an end tag (such as </author>), and
24 information between the two tags (which is referred to as the content of the objects). An
25 object may include a name-value pair (referred to as an attribute) related by an equal sign

1 that modifies certain features of the object (such as MONTH = "May"). The objects in
2 an XML document have a hierarchical relationship to each other that can be represented
3 as a tree, such as the tree shown in Fig. 2. The objects in the tree are also commonly
4 referred to as "nodes." A so-called XML schema provides a formal specification that
5 defines the types of objects and the organization of objects that should appear in an XML
6 document in order for that document to be considered so-called well formed.

7 Fig. 6 shows an exemplary schema 600 for defining the selection and organization
8 of information in templates. It includes hierarchical arrangement of elements (or objects)
9 (including elements 602, 604, 606, 608, 610, 612, 614, and 616). The elements are
10 denoted by the letter "E" in Fig. 6. Selected elements also include one or more attributes
11 associated therewith. The attributes are represented by the letter "A" in Fig. 6. More
12 specifically, Fig. 6 represents an XML schema 600 for a group of machines. Group
13 templates are a superset of machine templates. Each of the objects in the schema 600 is
14 identified in more detail below.

15 The root-level group element 602 represents a sequence of groups contained in
16 the template. The group element 602 has attributes of: Name (representing the name of
17 the group); Description (representing the description for the group); and DirName
18 (representing the subdirectory that contains APF files for the group). The default element
19 604 contains a list of parameters that are used as defaults for the group. The parameter
20 element 606 is used to define a parameter of an object. The parameter element 606
21 includes the attributes of: Name (representing the name of the parameter); Value
22 (representing the value of the parameter); Description (representing a description of the
23 parameter); Flags (representing flags associated with the parameter); Type (represent the
24 type of the parameter); and Display_Name (representing the displayed name associated
25 with the parameter).

1 The machine element 608 is used to define an instance of a machine. The
2 machine element 608 includes the following attributes: Name (representing the name of
3 the machine); Description (representing a description of the machine); UniqueID
4 (representing a unique ID assigned to the machine); Vendor (representing a vendor
5 associated with the machine; Model (representing a model associated with the machine);
6 OperatingSystem (representing an operating system associated with the machine); and
7 MachineFunction (representing a machine function associated with the machine). The
8 values element 610 contains a list of parameter values that are set on the machine. The
9 parameter object 612 specifies parameters associated with the machine. The following
10 attributes are associated with the parameter object 612: Name (referring to the name of
11 the parameter); Description (referring to the description of the parameter); and Value
12 (referring to the Value assigned to the parameter). The stage element 614 references a
13 particular stage. It has a stage attribute that identifies the name of the particular stage.
14 Finally, the PhaseList element 616 is used to set the phase list for a specific stage. It
15 includes the PhaseLink attribute that identifies the name of phase link that references a
16 phase list associated with the stage.

17 Fig. 7 shows an exemplary XML template 700 for a machine. As shown there,
18 this template 700 includes elements having attributes that correspond to the schema 600
19 shown in Fig. 6. Note that, in this template 700, the name of two phase lists are provided.
20 The specific phases in these phase lists and the associated scripts can be recovered from
21 the library 124.

22 Generally, references are used extensively within templates to reference objects in
23 the database 120. One format for referencing a desired object or parameter uses
24 expression references. The use of expression references was described above. Another
25 format uses so-called ID references. An ID reference represents a reference to a specific

1 Owner and Owner ID of the object. For example, the following ID expression references
2 an object with owner “MS” and ID88: “MS.88.”

3 Another way of referencing different parts of an XML template is by using XPath.
4 XPath is a general-purpose query notation for addressing and filtering the elements and
5 text of XML documents. XPath expressions can address parts of an XML document, and
6 can manipulate strings, numbers, and Booleans, etc. Using XPath, a value of a parameter
7 can be defined as an XPath query to another value in the template.

8 A template such as template 700 shown in Fig. 7 can be exported from one site to
9 another. One technique for performing this task is via the user interface 136 provided by
10 the build management logic 116. For example, in one exemplary implementation, using a
11 mouse device, a user can point to a particular machine object in the hierarchical tree of
12 objects associated with a particular site. Right clicking on the mouse device will provide
13 a menu that gives the user the option of exporting the XML template corresponding to the
14 designated machine to an identified destination, such as another site.

15 Fig. 8 provides an overview 800 that illustrates how library information 802 and
16 template information 804 combine to provide files 806 from which a specific machine
17 can be built. This procedure can come into play, for instance, when a template is
18 imported from another site and the user wishes to use this template to add a machine to
19 their build site. The library information 802 is obtained from the library 124 of Fig. 1.
20 The template information 804 is obtained from template storage 117 in Fig. 1. More
21 specifically, as described above, the library 124 stores a collection of objects that define
22 building blocks that a user can string together to provide configuration instructions for a
23 specific collection of machines (or other resources). The objects stored in the library 124
24 are maintained in a relatively general form. On the other hand, the template storage 117
25 provides information in a particular format that defines how to apply the information in

1 the library 124 to configure a specific set of machines at a site. The template information
2 804 may include, for instance, a number of parameters that are unique to the particular
3 machine being built. The template information 804 may also include references to more
4 general configuration objects present in library 124. When combined together, the
5 template information 804 and the library information 802 enable the build management
6 logic 116 to construct the files 806 to be used to build a machine or group of machines.

7 Another way to transfer information between different locations is through the
8 packaging functionality provided by the build management logic 116. More specifically,
9 a package refers to a self-contained collection of files that contain information extracted
10 from the database 120. More specifically, a package will contain all library information,
11 referenced files and other information required to completely define a design
12 environment for generating configuration instructions at a site.

13 Generally, a user may create a package by creating a package object (referring to
14 a root object that identifies the package), adding references to other objects to be included
15 in the package (such as phase lists, vendors, models, machine functions, operating
16 systems, and machine groups), adding parameters appropriate to the package, and then
17 generating the thus-defined package. A created package has the following attributes:
18 Name (representing the name of the package); Description (referring to a textual
19 description of the package); author (representing a textual description of the package's
20 author), and version (representing a version number assigned to the package, which is
21 updated upon each subsequent modification of the package). The inclusion of objects in
22 the package may also explicitly include other objects that are referenced.

23 As discussed above, a package can be created at one site (e.g., site X 102) and
24 transferred to another site (e.g., site Y 104). Further, a package at another site (e.g., site
25 Y 104) can be downloaded to a local site (e.g., site X 102). Further, a package can be

1 downloaded from the head-end site 166. Further still, a package can be uploaded to the
2 head-end site 166. The database 170 used to store templates at the head-end site 166 can
3 provide unified storage at a particular location, or can be implemented as a distributed
4 resource.

5 When the user has defined all objects and parameters to be used in configuring a
6 specific machine, then the user can generate the actual configuration instructions used to
7 build the machine. This can be performed in different ways. In one technique, the user,
8 using a mouse device, can point to a particular machine located in a hierarchy of objects
9 presented on the user interface 136 of the display device 134. Right-clicking on the
10 machine prompts the display of a menu. One item of the menu gives the user the option
11 to generate the configuration instructions for the identified machine.

12 Fig. 9 shows a simplified example of part of a file 900 that provides configuration
13 instructions. As indicated there, the file 900 includes information 902 that references
14 various stages involved in configuring the identified machine. For instance, the file 900
15 provides a list 904 of phases used to execute the PreOSStage. An exemplary phase 906
16 in this list (DelAllParts) points to more detailed instructions 908 for executing this phase.
17 Similarly, an exemplary phase 910 in the PostOSStage stage points to more detailed
18 instructions 912 for executing this stage. The thus-created .INI file is transferred to the
19 deployment module 146, and subsequently to the machines (152, 154, 156) to be used in
20 building these machines.

21

22 **B. User Interface (UI) Functionality**

23 Figs. 10-29 show different presentations provided by the system 100 of Fig. 1
24 used in performing a variety of tasks. More specifically, these presentations are
25 displayed on the user interface (UI) 136 of the display device 134. The layouts of these

1 UI presentations are exemplary; other implementations of the principles described here
2 can employ different kinds of UI presentations. Each of the UI presentations in Figs. 10-
3 29 will be described in sequence in the following explanation.

4

5 *Machine Definition UI Functionality*

6 The database 120 may initially contain no configuration-related information. A
7 user can provide such configuration-related information by importing it from another site
8 or from the head-end site 166. This task is performed by specifying the location of the
9 package and then importing the package from the identified address.

10 The imported package can contain templates for setting up a group of machines at
11 a particular site. The build management logic 116 allows users to fill in this template by
12 presenting a series of configuration guidance panels to the user (such as the well-known
13 wizard-type instruction panels). Figs. 10-13 provide UI presentations that show an
14 exemplary series of such panels.

15 To begin with, Fig. 10 shows a configuration guidance panel 1002 for specifying
16 a group type. The group type identifies a category of machine grouping to be built at a
17 particular site, typically associated with a particular machine group role. Fig. 10
18 specifically presents a subpanel that presents a collection 1004 of group types from which
19 the user can select. In this example, the user has selected a web hosting group type entry
20 1006. A subsequent configuration guidance panel (not shown) can allow the user to enter
21 information regarding the defaults applicable to the selected group.

22 Fig. 11 shows another configuration guidance panel 1102 for specifying the
23 machines that should be present in the group. More specifically, a left panel presents a
24 list 1104 of machines within the group. A user can check the boxes corresponding to
25 machines that should be present in the group. In this example, the user has selected all of

1 the machines. A right panel presents entry fields 1106 that allow the user to specify
2 details regarding the selected machines, such as: Name (representing the name of the new
3 machine); a Unique ID (representing a unique ID assigned to the machine); Operating
4 System (representing the operating system to be installed on the new machine); Vendor
5 (representing the manufacturer or provider of the new machine); Model (representing the
6 model of the new machine); and Function (representing the function or role assigned to
7 the new machine). In the present example, the user is currently specifying the properties
8 of a machine AD2 1108, highlighted in the right panel 1104. After specifying a group of
9 machines in the above-identified manner, the properties of the group can be displayed.
10 For instance, the build management logic 116 can present a hierarchical site tree that
11 identifies objects associated with the group of machines. The user can add additional
12 groups of machines in the above-described manner by pointing to a site object in the site
13 tree and then right clicking on the mouse device.

14 Fig. 12 shows another configuration guidance panel 1202 for specifying the
15 properties of a machine. Generation of this panel can be initiated by pointing to an
16 appropriate location in the site tree (such as a selected group within the site tree) and then
17 right clicking on the mouse device. This configuration guidance panel 1202 prompts the
18 user to enter a collection 1204 of information regarding the new machine, including:
19 Name, Description; Unique ID; INI File Name (representing the name of the .INI file
20 created for the new machine); Boot Type (representing the type boot device associated
21 with the new machine); Operating System; Vendor; Model; and Function. A subset of
22 the fields for specifying information within the collection 1204 can include respective
23 drop-down selection boxes for displaying a group of options available to the user in
24 defining the new machine. For instance, drop-down selection box 1206 shows the
25

1 following exemplary functions: Base Windows 2000 Server; DNS; Domain Controller;
2 and Web Server.

3 The configuration guidance panel 1302 shown in Fig. 13 includes a collection
4 1304 of drop-down boxes through which the user can define phase lists associated with
5 different stages used to configure the new machine. For instance, the exemplary drop-
6 down box 1306 allows the user to select between two specified phase lists. With respect
7 to Figs. 10-13, the configuration guidance procedure can dynamically change the options
8 presented to the users in the drop-down boxes depending selections made in prior drop-
9 down boxes.

10 Another way to create groups of machines or individual machines once a site tree
11 has been established is to use cut and paste and/or drag and drop functionality. In this
12 technique, a user can “cut” configuration objects from one part of the hierarchical tree,
13 and “paste” them to another part of the tree, or can use drag and drop functionality to
14 perform the same operation.

15

16 *Site Tree UI Functionality*

17 Fig. 14 is used as a vehicle for introducing the basic exemplary layout of the UI
18 presentation 136. The presentation includes a tree pane 138 that provides a hierarchical
19 tree listing of objects associated with the database 120. More specifically, a first portion
20 1402 of the tree pane 138 provides a tree listing of objects pertaining to machines
21 associated with a particular site (where the site is arbitrarily named “BRUCECH” in this
22 example) (corresponding to site tree 126 of Fig. 1). A second portion 1404 of the tree
23 pane 138 provides a tree listing of objects associated with packages (corresponding to
24 packages 120 of Fig. 1). The second portion 1404 is not expanded to reveal its contents
25 in Fig. 14, but in later figures the tree will be exploded to reveal its contents. A third

portion 1406 of the tree pane 138 provides a tree listing of objects associated with the library 124 of Fig. 1. A bottommost object 1408 in the library provides information regarding permissible types of parameters used in the build management logic 116. Types of parameters can include: boot type; directory type; file type; integer type; IP address auto type; IP/Mask static type; on/off binary type; password type, and so on. The tree display pane 138 includes a vertical scroll bar 1410 used to move the displayed tree up and down to reveal different parts of it.

A parameter display pane 140 provides information regarding parameters associated with different objects in the tree listing shown in the tree pane 138. In this case, the exemplary machine object WEB1 1412 in the tree pane 138 is highlighted. This prompts the build management logic 116 to display a collection 1414 of parameters associated with this machine at the machine level. More specifically, the parameter display pane 140 provides a first field that identifies the name of the parameter, a second field that defines type of the parameter, a third field that defines the value assigned to the parameter (if it can be determined), and a fourth field that defines the level in the hierarchy at which the parameter was initially created.

A properties display pane 142 provides information regarding the properties of an object that is highlighted in the tree pane display 138 – in this case, WEB1 machine 1412. Hence, the properties display pane 142 provides machine properties associated with object 1412. A specific collection 1416 of properties includes: Name, Description, Unique ID, INI File Name, Vendor, Model, Operating System, and Function assigned to the machine. An apply button 1418 records the selections made by a user via the properties display pane 142. A tool bar 1420 located at the top of the window-type UI object shown in Fig. 14 can include a collection of icons (not shown) for initiating other actions with respect to the creation and management of configuration instructions.

1 Fig. 15 shows a UI presentation for displaying properties associated with
2 parameters. In the specific example of Fig. 15, the user has highlighted one of the
3 parameters in the parameter display pane 140, namely parameter 1502. Parameter 1502,
4 in turn, is associated with highlighted machine object WEB2 1504 in the tree display
5 pane 138. By virtue of the highlighting of a parameter in the parameter display pane 140,
6 the build management logic 144 displays properties associated with the identified
7 parameter in the property display pane 142. For instance, the property display pane 142
8 represents a collection 1506 of properties including: Name information (representing the
9 name assigned to parameter 1502); Type (representing the type of the parameter 1502);
10 Value 1508 (representing the value assigned to the identified parameter 1502); and Flags
11 1510, representing flags associated with the parameter 1502). In this case, the Value
12 1508 of the parameter 1502 is assigned the IP address “10.18.1.”

13 The Flags 1510 assigned to the parameter 1502 include a collection of fields that
14 can be toggled on or off (indicated by a check mark, or the absence of a check mark,
15 placed in the boxes associated with the flags). More specifically, a “quote” flags prompts
16 the build management logic to add quotes around the parameter value when it appears in
17 the .INI file. A “hide name” flag prompts the build management logic 116 to cause only
18 data to be written (instead of the typical format of “name = value”). A “suppress” flag
19 prompts the build management logic 116 to not write the parameter value to the .INI file.
20 An “advanced” flag prompts the build management logic 116 to display advanced
21 properties associated with the parameter, such as an identification of the owner of the
22 parameter. A “machine specific” field prompts the build management logic 116 to
23 display the parameter at the machine level and reset the parameter when a copy of the
24 machine is made. Finally, a “mandatory” flag prompts the build management logic 116
25 to highlight the parameter when it appears in the parameter display field 140 when its

1 value is mandatory and is also left blank. For instance, parameter 1512 shown in the
2 parameter display field 140 is displayed in red to indicate that it is mandatory and its
3 value has been omitted.

4 Caption 1514 provides information regarding the owner the of highlighted
5 parameter 1502, as well as the version of the highlighted parameter.

6 Whereas Fig. 15 presented a case where the highlighted parameter was assigned a
7 fixed value (“10.18.1”), Fig. 16 presents a case where a highlighted parameter 1602 is
8 assigned a variable value 1604 (i.e., “=Default.DistUsr”). The parameter 1602 is
9 associated with a common stage 1606 in the configuration of the WEB2 machine. The
10 user can prompt the build management logic 116 to resolve the value of the variable
11 value 1604. The UI presentation responds by displaying the resolved value in window
12 1608 (identifying that the parameter DistUser is resolved to have a value of:
13 “administrator”).

14 Fig. 17 shows an “advanced view” feature provided by the build management
15 logic 116. In one implementation, this feature can be implemented by activating a
16 corresponding “advanced view” menu item under the view menu. The advanced view
17 feature prompts the build management logic 116 to display additional information
18 regarding the objects and properties shown in the UI presentation 128. For instance, the
19 advanced view feature can prompt the build management logic 116 to display
20 information regarding the individual who has rights to change the objects and properties.
21 For example, added textual information indicates that the user BRUCECH has rights to
22 modify the common stage object 1702, and the parameter 1704. The individual who has
23 rights to modify these features typically corresponds to person who originally created or
24 specified these features.

1 *Library UI Functionality*

2 The UI presentation 128 provided by the build management logic 116 also allows
3 a user to perform various actions that target the library 124. As discussed in Section A,
4 the library 124 stores the universe of objects available to the user when defining a
5 particular object configuration for a site. Fig. 18 is used as an introductory vehicle to
6 illustrate selected features of the library-based UI functionality provided by the build
7 management logic 116. To begin with, the tree display pane 138 shown in Fig. 18
8 includes a breakdown of categories of objects available in the library 124. As was
9 discussed in connection with Fig. 4, the library includes vendor information 1802,
10 operating system information 1804, machine function information 1806, stage, phase, and
11 phase list information 1808, and parameter type information 1810. The common stage
12 1812 is highlighted in the tree display pane 138, which prompts the build management
13 logic 116 to show the parameters 1814 associated with this stage in the parameter
14 display pane 140. Various properties 1816 of the common stage are presented in the
15 properties display pane 142.

16 Fig. 19 shows the display of information regarding a phase list within the database
17 120. More specifically, an identified phase list 1902 in the library tree identifies a
18 collection 1904 of items that specify various kinds of information, such as links to
19 phases, and links to other phase lists. For example, item 1906 represents a link to another
20 phase list. Item 1908 represents a link to another phase. Information 1910 in the
21 property display pane 142 conveys details regarding properties associated with the
22 identified phase list 1902.

23 In Fig. 20, the user has selected a particular phase link 2002 within the phase list
24 2004. The parameters 2006 displayed in the parameter display pane 140 pertain to the
25 phase that the phase link is referencing. The user can override the parameter values at

1 this level. Such an override will only affect the parameter value in the context of the
2 phase list 2004. The properties display pane 142 presents information regarding the
3 identified phase link 2002.

4 To add a phase to a phase list, the user can copy the desired phase from another
5 source. The user can then select the phase list that is to receive the desired phase, and
6 then point to the phase link below where the user wishes the new phase to be inserted.
7 Then the user can paste the new phase into the specified location, prompting the build
8 management logic 116 to define a new phase link.

9 As shown in Fig. 21, a phase list can also have a special type of link called a
10 reference. A reference points at a parameter using the convention “category.parameter.”
11 The referenced category parameter, in turn, should point to a phase or a phase list. For
12 instance, the reference 2102 in Fig. 21 points to the phase list “PostOSPhaseList” in the
13 category “OperatingSystem.” This feature allows a user to define a phase list that varies
14 depending on the operating system (or other category) chosen.

15 Fig 22 shows a procedure for entering a phase to the library of phases. To
16 perform this task, the user can right click the mouse device on the category “phases.any”
17 2202 in the library tree within the tree display field 138. This prompts the build
18 management logic 116 to display the script selection panel 2204 that identifies a list of
19 available scripts 2206. Fig. 22 indicates that the user has selected script 2208 within the
20 list of available scripts 2206. The user can also specify a new name for the selected
21 script. Fig. 22 shows the inclusion of a phase link 2210 for the new pane within the tree
22 display pane 138.

23 Fig. 23 shows a shorthand technique for adding phase parameters to a phase. In
24 the case of Fig. 23, the technique is used to add parameters to a phase represented by
25 phase link 2302. In this technique, the user retrieves and displays an already created

1 manifest file 2304. The manifest file 2304 contains a collection 2306 of parameters in its
2 body. The user can associate this collection 2306 of parameters with the phase pointed to
3 by phase link 2302 by cutting and pasting the collection 2306 of parameters from the
4 manifest file to the parameter display pane 140. Another technique for adding new
5 parameters is to right click on the mouse device to command the build management logic
6 116 to add a new parameter as specified by the user.

7 Fig. 24 shows a technique for adding data types and flags to a specified parameter
8 2402 associated with phase link 2404. A drop-down menu 2406 within property display
9 pane 142 is used to specify the parameter type associated with the parameter 2402, while
10 field 2408, containing a collection of check boxes, is used to specify the flags associated
11 with the parameter 2402. Validation rules can enforce the specification of types made via
12 the drop-down menu 2406. That is, the build management logic 116 can be configured to
13 generate an error message if the user attempts to enter information that does not conform
14 to the data types specified via the drop-down menu 2406. This task is performed by
15 comparing the user's input with predetermined validation criteria.

16 More generally, in another implementation, the build management logic 116 can
17 include functionality that enables the user to specifically tailor the validation rules that
18 will be invoked upon the entry of data into the build management logic 116. For
19 instance, the build management logic 116 can include an interface panel (not shown)
20 which allows the users to add, modify, and delete validation rules, and to fully define the
21 behavior of the validation rules (e.g., when they are invoked, in what context they are
22 invoked, the messages that will be displayed in response thereto, and so on).

23 Fig. 25 indicates that the user has selected the data type directory for the
24 parameter 2402 associated with phase 2404. This prompts the build management logic
25 116 to generate a directory value attribute display panel 2502. This panel 2502 specifies

1 the information required to get the install files for this phase 2404. For instance, panel
2 2502 can specify the location of the install files by providing local path information or
3 URL information in selection fields 2504 and 2506, respectively. The build management
4 logic 116 can be configured to check whether the specified file exists upon generating the
5 manifest files, rather than during the actual configuration procedure.

6

7 *UI Features Pertaining to the Generation of Output Files and Packages*

8 After the user is satisfied with the object-specification of a site, the user can
9 export one or more XML templates that describe configuration instructions used to build
10 respective machines. Fig. 26 shows an exemplary procedure for exporting an XML
11 template for an identified WEB2 machine 2602. The procedure involves pointing to the
12 machine object 2602 in the tree display pane 138 with a mouse device, and then right-
13 clicking on the machine object. This prompts the build management logic 116 to display
14 a task menu 2604. Selection of an “All Tasks” item 1506 in the task menu 2604 prompts
15 the build management logic 116 to generate another task menu 2608. The other task
16 menu 2608 has an item 2610 that commands the build management logic 116 to export
17 the machine template for the WEB2 machine object 2602. An exemplary exported XML
18 template is illustrated in Fig. 8. The user can export an XML template for an entire group
19 of machines in the above-described manner by right clicking on a group of machines in
20 the site tree.

21 Instead of exporting an XML template for the machine object 2602, the user may
22 want to generate an .INI file that defines the configuration instructions for this machine
23 object 2602. Fig. 27 provides a technique for performing this task. Again, the user can
24 initiate this task by right-clicking on the WEB2 machine object 2602, which prompts the
25 display of the menu 2604. Clicking on a “Generate Files” option 2702 in the menu 2604

1 prompts the generation and display of the .INI file 1604 for the machine object 2602. As
2 shown in Fig. 27, the build management logic 116 displays the .INI file 2604 using a text
3 editor 2706. If the user is satisfied with the .INI file, the user can then command the
4 system 100 shown in Fig. 1 to configure the physical machine corresponding to the
5 WEB2 machine object 2602 using the .INI file.

6 Once the configuration operation is underway, the user can use the build
7 management logic 116 to also report the status of configuration. This feature is
8 illustrated in Fig. 28. As indicated there, the user has initiated the configuration of a
9 machine corresponding to machine objects in group 2802. The user can then determine
10 the status of the configuration for each of the machine objects in the group 2802. A
11 status inquiry can be initiated through various mechanisms, such as by clicking on a
12 corresponding status-related icon in the tool bar (not shown).

13 Status information can be conveyed via the status display pane 2804 shown in
14 Fig. 17. This status display pane 2804 includes numerical status information in field
15 2806 and graphical status information in field 1708. More specifically, the field 2808
16 conveys the level of completion of configuration tasks in bar chart format for the
17 respective machines listed in field 2810.

18 Finally, Fig. 29 shows the management of packages using the UI presentation
19 128. Namely, the tree display pane 138 shows a collection of packages 2902. A specific
20 package 2904 is selected and its contexts expanded for viewing. The parameter display
21 pane 140 shows parameters associated with the selected package 2904. The property
22 display pane 142 shows properties associated with the selected package 2904. A new
23 package can be created by activating a corresponding menu item that initiates the creation
24 of the package. The user can then drag and drop phase lists and other objects from the
25 library into the newly created package. A final package can contain phase lists, phases,

1 and associated script files, etc. in a single package that can easily be transferred between
2 users at different sites.

3

4 *UI Features Pertaining to Table View Functionality*

5 The above-described Figs. 14-29 use a so-called tree view to present information
6 regarding objects in the tree pane 138. The tree view allows the user to display a
7 hierarchical list of machines within a group by double left clicking on the group in the
8 site tree (e.g., the site tree 1402 in Fig. 14). The user can then select an object within the
9 expanded hierarchical list and view the properties of the object in the manner described
10 above.

11 In contrast, Figs. 30-33 use a so-called table view to present information
12 regarding objects in the tree pane 138. As shown in Fig. 30, a table view 3000 provides
13 information regarding machines in a group of machines using a table format. For
14 instance, in Fig. 30, the user has selected the machine group “Web Host Test Lab
15 Sample” 3002. The table view 3000 presents a table that identifies the machines in the
16 group 3002 and also provides information regarding the machines, such as Name,
17 UniqueID, IniFile, Boot.Management, Common.Sync, InstOSStage.Para,
18 InstOSStage.Para, etc. The properties display pane 142 presents information regarding
19 one or more selected machines in the table view 3000. For instance, the user in Fig. 30
20 has selected object 3004 in the table view 3000. The properties display pane 142
21 responds by providing information regarding the properties of this object 3004. More
22 specifically, the properties display pane 142 allows the user to view: a) general
23 information regarding the selected object; b) information regarding the phase list
24 properties associated with the selected object; and c) information regarding the parameter
25 properties associated with the selected object. In the case of Fig. 30, the user has selected

1 a pane that provides general information 3006 regarding the selected object. The
2 properties display pane 142 allows the user to edit and modify the properties of the
3 selected machine – in this case, machine 3002.

4 To switch between tree view and table view, the user can right click on a selected
5 group in the tree pane 138. This causes the presentation of a menu that gives the user the
6 option of selected either the tree view or the table view. One reason that the user may
7 want to select the table view is to edit groups of machines in an efficient en bloc manner.
8 This editing technique is discussed further with respect to Figs. 31-33.

9 To begin with, in Fig. 31, the user has selected a group of machines 3102 in the
10 table view 3000. This can be performed by left clicking and dragging on the row
11 selection field 3104 using the mouse device. In the properties display pane 142, the user
12 has opted to view the properties of the phase lists used in configuring the selected group
13 of machines 3102. That is, a collection of fields 3106 shows information regarding the
14 phase lists used in the PreOSStage, InstOSStage, and PostOSstage stages when
15 configuring the selected group of machines 3102. Since multiple machines 3102 have
16 been selected, there is a chance that different machines will use different phase lists. In
17 this case, the build management logic 116 present blanks fields for any dissimilar phase
18 lists in the collection 3106. However, in the example of Fig. 31, all of the machines in
19 the selected group of machines 3102 use the same phase lists. The collection of fields
20 3106 therefore shows the common phase list information for the selected group of
21 machines 3102. Any changes made to the collection of fields 3106 will modify the phase
22 lists associated with all of the machines in the selected group of machines 3102 (that is,
23 upon the user pressing the Apply control button). Accordingly, this editing technique
24 provides a quick and efficient mechanism for changing common properties associated
25 with multiple machines.

1 The build management logic 116 also gives the user the option of exporting the
2 selected group of machines 3102 to a CVS file (i.e., a Comma-Separated Values file) or
3 other kind of file. Information in this file can then be manipulating using conventional
4 spreadsheet applications, such as the Excel software product produced by Microsoft
5 Corporation of Redmond, Washington. The build management logic 116 also gives the
6 user the option of importing CVS files. Thus, the user can export a group of machines,
7 edit the machines using a spreadsheet application, and then import the edited machines
8 back into the build management logic 116 for viewing on its UI.

9 Fig. 32 shows the above-described case where the user has selected a group of
10 machines 3202 that employ dissimilar phase lists. In this case, the collection of fields
11 3204 in the properties display pane 142 will present blank fields.

12 Fig. 33 shows the case where the user has opted to view the parameters associated
13 with a selected group of machines 3302. In this case, the build management logic 116
14 provides a table 3304 that identifies the parameters associated with the selected group of
15 machines 3302. By default, the build management logic 116 shows machine-specific
16 parameters. Additional parameters can be shown by right clicking and selecting an “add
17 parameter” option presented to the user. Again, if any of the parameters differ within the
18 selected group of machines 3302, then the build management logic 116 displays a blank
19 field for these parameters. A change to any of the parameter values shown in the table
20 3304 will affect all of the machines in the selected group of machines 3302.

21

22 **C. Exemplary Method of Operation**

23 Fig. 34 shows an overview of an exemplary procedure 3400 for creating and
24 deploying configuration files using the system of Fig. 1. Step 3402 entails specifying the
25 objects that govern the generation of the configuration files using the build management

1 logic 116. The objects are arranged to form a hierarchical “site tree” shown in the UI
2 presentation figures (i.e., Figs. 10-29). The site tree can be produced using various
3 techniques. Step 3404 entails building the tree from scratch by successively defining the
4 properties of machines which make up the site tree based on a template. Step 3406
5 entails importing a site tree from another site or the central site 166. Step 3408 entails
6 modifying a pre-existing site tree (if the user has permission to make such a modification)
7 to produce a new site tree.

8 Step 3410 entails defining the values for any parameters that are mandatory yet
9 remain unspecified. In one implementation, the build management logic 116 highlights
10 critical unspecified parameters in the parameter display pane 140 of the UI presentation
11 128.

12 Step 3412 entails generating the manifest files based on the site tree defined in
13 step 3402. Step 3414 entails actually building the machines based on the manifest files
14 created in step 3412.

15 Step 3416 entails optionally exporting XML templates for a machine or group of
16 machines to a destination site. Step 3416 can also involve creating a package and
17 transferring the package to the destination site.

18 Fig. 35 shows an exemplary procedure 3500 whereby parameters associated with
19 objects in a configuration hierarchy can inherit values from objects located higher in the
20 hierarchy. Step 3502 involves adding a configuration object to a site tree, where the
21 object has at least one parameter having an undetermined value. Step 3504 involves
22 resolving the undetermined value. Step 3506 involves generating configuration files
23 from the site tree once the value has been resolved. (However, in one implementation,
24 the build management logic 116 does not perform step 3504 until the user enters a
25 request to generate the configuration files.)

1 The step 3504 of resolving the value can involve performing the procedure shown
2 in the right hand side of Fig. 35. In that procedure, step 3508 involves looking to a next
3 higher node in the site tree to determine whether that node can resolve the undetermined
4 value. If the value is resolved, as assessed in step 3510, then the procedure in step 3504
5 terminates. If the value is not resolved, then the flow advances to step 3512, which
6 determines whether the top of the site tree has been reached without determination of the
7 unspecified value. If this is the case, then the build management logic 116 can abandon
8 the procedure shown in step 3504 and so notify the user (e.g., by highlighting the
9 parameter having the missing value in the parameter display pane 140 of the UI interface
10 128). If the highest node has not been reached, the build management logic 116 repeats
11 the above-described steps with respect to the next higher node in the site tree. In this
12 manner, an initially undetermined value associated with an object in the site tree can
13 inherit the value defined in a higher level of the site tree. This makes it easy to modify
14 the site tree by adding and deleting objects in the site tree without having to perform
15 pain-staking and ad hoc revision of the tree as a whole.

16 Other kinds of referencing strategies can be used to supplement the inheritance
17 strategy described in procedure 3500.

18 Fig. 36 shows an exemplary procedure 3600 for building a machine by taking into
19 account interdependencies involved in configuring multiple different machines. The left
20 side 3602 of the figure shows operations performed locally at a machine being built
21 (referred to as the “specified machine”). The right side 3604 of the figure shows
22 operations performed at the deployment module 146 of Fig. 1.

23 Beginning with the left side 3602, step 3606 entails initiating the configuration of
24 the specified machine based on previously defined manifest files for that machine. Step
25 3608 entails performing a configuration action specified in the manifest files. For

1 instance, this step may involve executing a script associated with a phase specified in the
2 manifest files, or performing part of a script associated with the phase. Step 3610
3 involves determining whether it is time to report the specified machine's progress in
4 building the specified machine to the deployment module 146. For instance, the
5 specified machine may report its progress at predetermined intervals, or after
6 predetermined events have taken place, or in response to an express polling request by
7 the deployment module 146. If step 3610 is answered in the affirmative, then step 3612
8 involves sending a report to the deployment module 146. Following steps 3610 and
9 3612, the specified machine determines whether it has reached a point in its configuration
10 such that it needs to take into account dependencies between its own configuration
11 progress and the configuration progress of other machines at the site. For instance, the
12 specified machine may require a value that is furnished by other machines only when the
13 other machines reach a predetermined point in their own configuration. Many other
14 factors may create configuration interdependencies among multiple machines. Whatever
15 the cause, at this juncture, step 3614 entails determining whether the specified machine
16 needs to check the progress of the configuration of other machines before proceeding
17 further with its own configuration. If this is so, step 3616 entails determining the status
18 of other machines by querying the master dependency log 150. Once the specified
19 machine is granted approval to proceed with its configuration, the flow advances to step
20 3618, in which it is determined whether the configuration of the specified machine has
21 been completed. If this is the case, then, in step 3620, the configuration terminates. If
22 the configuration has not finished, then the specified machine repeats the above-identified
23 steps, that is, by performing the configuration of the machine in piecemeal fashion,
24 stopping whenever necessary such that the other machines can catch up with it if need be.
25

1 Now referring to the right side 3604 of the figure, step 3622 generally represents
2 the logging of progress information obtained from different machines within the master
3 dependency log 150 of the deployment server 146. More specifically, as mentioned
4 above, each machine periodically sends status reports to the master dependency log 150
5 in the deployment module 146. The status reports reflect each machine's progress in
6 performing its allocated configuration tasks.

7 A machine being configured can periodically poll the information in the master
8 dependency log 150, or can access the master dependency log 150 upon the occurrence of
9 specific events (such as the machine reaching a predetermined point in its allocated
10 configuration tasks). Information retrieved from the master dependency log 150 specifies
11 how far other machines have advanced in their own configuration. This information, in
12 turn, allows the machine to determine whether it can advance to a subsequent
13 configuration task.

14 In the above discussion, a machine being configured played an active role in
15 polling the master dependency log 150, and then analyzing information obtained
16 therefrom. However, may other strategies can be used to achieve the same monitoring
17 function. For instance, the deployment module 146 can play an active role by
18 periodically sending reports to each machine being configured which alerts each machine
19 as to the configuration status of other machines. Alternatively, the deployment module
20 146 can include logic for resolving dependency synchronization issues, rather than
21 allocating this determination to each individual machine being configured.

22
23
24
25

D. Exemplary Computer Environment for Implementing the Configuration Management and Definition Module

Fig. 37 shows an exemplary computer environment 3700 that can be used to implement any of the processing devices shown in Fig. 1, such as the configuration management and definition module 114 or the deployment module 146. The computing environment 3700 includes the general purpose computer 3702 and display device 3704 (which can correspond to the display device 134 shown in Fig. 1). However, the computing environment 3700 can include other kinds of computer and network architectures. For example, although not shown, the computer environment 3700 can include hand-held or laptop devices, set top boxes, programmable consumer electronics, mainframe computers, gaming consoles, etc. Further, Fig. 37 shows objects of the computer environment 3700 grouped together to facilitate discussion. However, the computing environment 3700 can employ a distributed processing configuration. In a distributed computing environment, computing resources can be physically dispersed throughout the environment.

Exemplary computer 3702 includes one or more processors or processing units 3706, a system memory 3708, and a bus 3710. The bus 3710 connects various system components together. For instance, the bus 3710 connects the processor 3706 to the system memory 3708. The bus 3710 can be implemented using any kind of bus structure or combination of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. For example, such architectures can include an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

1 Computer 3702 can also include a variety of computer readable media, including
2 a variety of types of volatile and non-volatile media, each of which can be removable or
3 non-removable. For example, system memory 3708 includes computer readable media in
4 the form of volatile memory, such as random access memory (RAM) 3712, and non-
5 volatile memory, such as read only memory (ROM) 3713. ROM 3713 includes an
6 input/output system (BIOS) 3714 that contains the basic routines that help to transfer
7 information between objects within computer 3702, such as during start-up. RAM 3712
8 typically contains data and/or program modules in a form that can be quickly accessed by
9 processing unit 3706.

10 Other kinds of computer storage media include a hard disk drive 3714 for reading
11 from and writing to a non-removable, non-volatile magnetic media, a magnetic disk drive
12 3716 for reading from and writing to a removable, non-volatile magnetic disk 3718 (e.g.,
13 a “floppy disk”), and an optical disk drive 3720 for reading from and/or writing to a
14 removable, non-volatile optical disk 3722 such as a CD-ROM, DVD-ROM, or other
15 optical media. The hard disk drive 3714, magnetic disk drive 3716, and optical disk drive
16 3720 are each connected to the system bus 3710 by one or more data media interfaces
17 3724. Alternatively, the hard disk drive 3714, magnetic disk drive 3716, and optical disk
18 drive 3720 can be connected to the system bus 3710 by a SCSI interface (not shown), or
19 other coupling mechanism. Although not shown, the computer 3702 can include other
20 types of computer readable media, such as magnetic cassettes or other magnetic storage
21 devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical
22 storage, electrically erasable programmable read-only memory (EEPROM), etc.

23 Generally, the above-identified computer readable media provide non-volatile
24 storage of computer readable instructions, data structures, program modules, and other
25 data for use by computer 3702. For instance, the readable media can store the operating

1 system 3726, one or more application programs 3728 (such as the definition file module
2 118), other program modules 3730, and program data 3732.

3 The computer environment 3700 can include a variety of input devices. For
4 instance, the computer environment 3700 includes the keyboard 3734 and a pointing
5 device 3736 (e.g., a “mouse”) for entering commands and information into computer
6 3702. The computer environment 3700 can include other input devices (not illustrated),
7 such as a microphone, joystick, game pad, satellite dish, serial port, scanner, card reading
8 devices, digital or video camera, etc. Input/output interfaces 3738 couple the input
9 devices to the processing unit 3706. More generally, input devices can be coupled to the
10 computer 3702 through any kind of interface and bus structures, such as a parallel port,
11 serial port, game port, universal serial bus (USB) port, etc.

12 The computer environment 3700 also includes the display device 3704. A video
13 adapter 3740 couples the display device 3704 to the bus 3710. In addition to the display
14 device 3704, the computer environment 3700 can include other output peripheral devices,
15 such as speakers (not shown), a printer (not shown), etc.

16 Computer 3702 can operate in a networked environment using logical connections
17 to one or more remote computers, such as a remote computing device 3742. The remote
18 computing device 3742 can comprise any kind of computer equipment, including a
19 general purpose personal computer, portable computer, a server, a router, a network
20 computer, a peer device or other common network node, etc. Remote computing device
21 3742 can include all of the features discussed above with respect to computer 3702, or
22 some subset thereof.

23 Any type of network can be used to couple the computer 3702 with remote
24 computing device 3742, such as a local area network (LAN) 3744, or a wide area
25 network (WAN) 3746 (such as the Internet). When implemented in a LAN networking

1 environment, the computer 3702 connects to local network 3744 via a network interface
2 or adapter 3748. When implemented in a WAN networking environment, the computer
3 3702 can connect to the WAN 3746 via a modem 3750 or other connection strategy. The
4 modem 3750 can be located internal or external to computer 3702, and can be connected
5 to the bus 3710 via serial I/O interfaces 3752 other appropriate coupling mechanism.
6 Although not illustrated, the computing environment 3700 can provide wireless
7 communication functionality for connecting computer 3702 with remote computing
8 device 3742 (e.g., via modulated radio signals, modulated infrared signals, etc.).

9 In a networked environment, the computer 3702 can draw from program modules
10 stored in a remote memory storage device 3754. Generally, the depiction of program
11 modules as discrete blocks in Fig. 37 serves only to facilitate discussion; in actuality, the
12 programs modules can be distributed over the computing environment 3700, and this
13 distribution can change in a dynamic fashion as the modules are executed by the
14 processing unit 3706.

15 Wherever physically stored, one or more memory modules 3708, 3718, 3722, etc.
16 can be provided to store the build management logic 116 programming code.

17 The server 166 illustrated in Fig. 1 can include a similar architecture to that
18 described above, e.g., including one or more processing devices, one or more busses,
19 system memory, input/output interfaces, and so on.

20 The design concepts disclosed herein are applicable to other business and
21 technical environments. For instance, the design concepts are extendable to include
22 management and resource mapping for other resources, such as network device ports,
23 storage device resources, etc. In other application, the efficient transfer of XML template
24 documents can be used to implement a software vending or subscription system.

1 **E. Conclusion**

2 A technique for abstracting the task of generating configuration instructions was
3 described. The technique has the potential of expediting and facilitating the generation of
4 configuration instructions, and for facilitating the transfer of configuration instructions
5 between different users.

6 Although the invention has been described in language specific to structural
7 features and/or methodological acts, it is to be understood that the invention defined in
8 the appended claims is not necessarily limited to the specific features or acts described.
9 Rather, the specific features and acts are disclosed as exemplary forms of implementing
10 the claimed invention.

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25